What is an Implicit Neural Representation?

For any given discretely represented signal (e.g. audio as a waveform, images as grids of pixels, 3D shapes as voxel grids), there exists a function mapping it to a continuous representation that corresponds to the discrete representation at any point on its domain. An Implicit Neural Representation (INR) is an approximation of such a function achieved by fitting a multi-layer perceptron on the target signal.

INRs have uses derived from representing an input. For example, since they can be queried at any point of their domain, they are not bound to a specific resolution which makes possible data compression and super-resolution or signal denoising. INRs can also be used to reconstruct a signal from its derivatives, novel view synthesis and inverse problems. Two popular papers in the field are SIREN and NERF.

An issue that INRs experience is an inability to capture fine and rapidly changing patterns in data. This is called Spectral Bias and is addressed by using positional encodings or unorthodox activation functions, with supplementary techniques for improving their effectiveness.